

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (Currently Amended) A method of logging ~~updates and recovering from failure~~ ~~recovery~~ in a transaction-processing system having a main memory for storing a database, one or more persistent backup storage devices for storing a copy of the database in the main memory, and one or more persistent log storage devices for storing log records ~~for parallel logging and parallel recovery~~, the logging method comprising:

generating differential log records by applying a bit-wise exclusive-OR (XOR) operation between a before-update image and an after-update image, and

~~distributing the generated differential log records in parallel to said persistent log storage devices; and~~

~~recovering from a failure by replaying the differential records in an arbitrary order, which is independent of the order of generation of the log records, by using bit wise XOR operations~~

the recovery method comprising:

redoing updates of committed transactions by applying the bit-wise XOR operation between the differential log records read from said one or more persistent log

storage devices and the copy of the database read from said one or more persistent backup storage devices; and

undoing updates of uncommitted transactions by applying the bit-wise XOR operation between the differential log records read from said one or more persistent log storage devices and the copy of the database read from said one or more persistent backup storage devices,

wherein the log records are applied in a sequence independently from the order of log creation.

2. (Original) The method of Claim 1, wherein the database comprises a plurality of fixed-size pages.

3. (Currently Amended) The method of Claim 2, wherein said each log record has a log header comprising:

LSN (Log Sequence Number) for storing a log sequence;

TID(Transaction ID) for storing the identity of the transaction that created the log record;

Previous LSN for storing the identity of the most recently created log by the same transaction;

Type for storing the type of the log record;

Backup ID for storing the relation between the log record and the updated

page for use with fuzzy checkpointing;

Page ID for storing the identity of an updated page;

Offset for storing the starting offset of an updated area within the updated page; and

Size for storing the size of the updated area.

4. (Previously Presented) The method of Claim 1, further comprising:

checkpointing by occasionally writing the database in the main memory to said one or more persistent back storage devices.

5. (Previously Presented) The method of Claim 4, wherein checkpointing uses a transaction consistent checkpointing policy.

6. (Previously Presented) The method of Claim 4, wherein checkpointing uses an action consistent checkpointing policy.

7. (Previously Presented) The method of Claim 4, wherein checkpointing uses a fuzzy checkpointing policy.

8. (Currently Amended) The method of Claim 4, wherein ~~recovering the recovery method further~~ comprises:

loading the checkpointed database from said one or more persistent backup storage devices into the main memory database; and

loading the log records from said one or more persistent log storage devices into the main memory database in order to restore the main memory database to the most recent consistent state.

9. (Currently Amended) The method of Claim 8, wherein loading the checkpointed database is executed in parallel by partitioning ~~the backup data in said one or more backup storage devices~~.

10. (Currently Amended) The method of Claim 8, wherein ~~replaying the log records the recovery method~~ is done in two passes by separating a redoing pass and an undoing pass.

11. (Currently Amended) The method of Claim 10, wherein reading the log records and ~~replaying redoing/undoing~~ the log records are executed in a pipeline.

12. (Currently Amended) The method of Claim 10, wherein reading the log records is executed in parallel by partitioning the log records as well as ~~replaying redoing/undoing~~ the log records.

13. (Currently Amended) The method of Claim 12, wherein reading the log records and ~~replaying redoing/undoing~~ the log records are executed in a pipeline.

14. (Currently Amended) The method of Claim 8, wherein ~~replaying redoing/undoing~~ the log records is done in one pass.

15. (Currently Amended) The method of Claim 14, wherein reading the log records and ~~replaying redoing/undoing~~ the log records are executed in a pipeline.

16. (Currently Amended) The method of Claim 14, wherein reading the log records and ~~replaying redoing/undoing~~ the log records are executed in parallel by partitioning the log records.

17. (Currently Amended) The method of Claim 16, wherein reading the log records and ~~replaying redoing/undoing~~ the log records are executed in a pipeline.

18. (Previously Presented) The method of Claim 8, further comprising filling the main memory database with 0s in advance.

19. (Currently Amended) The method of Claim 18, wherein loading the checkpointed database comprises:

reading the checkpointed database from said one or more backup storage

devices; and

playing redoing/undoing the checkpointed database by applying the XOR operation between the checkpointed database and the main memory database.

20. (Currently Amended) The method of Claim 19, wherein reading the checkpointed database and playing redoing/undoing the checkpointed database are executed in a pipeline.

21. (Currently Amended) The method of Claim 19, wherein reading the checkpointed database is executed in parallel by partitioning the checkpointed database as well as playing redoing/undoing the checkpointed database.

22. (Currently Amended) The method of Claim 21, wherein reading the checkpointed database and playing redoing/undoing the checkpointed database are executed in a pipeline.

23. (Previously Presented) The method of Claim 19, wherein loading the checkpointed database and loading the log records are executed in parallel.

24. (Current Amended) A transaction processing system allowing logging updates and recovering recovery from failure, comprising :

a main memory for storing a database;

~~a plurality of one or more~~ persistent log storage devices for storing log records ~~for parallel logging and parallel recovery~~;

one or more persistent backup storage devices for storing a copy of the database in the main memory;

means for generating differential log records by applying [[a]] bit-wise exclusive-OR (XOR) operations between [[a]] before-update image images and [[an]] after-update [[image]] images;

~~means for distributing the generated differential log records in parallel to said persistent log storage devices~~; and

means for replaying the differential log records in an arbitrary order, independent of their generation order, by using the bit-wise XOR operations,

wherein said means for replaying further comprises means for redoing committed transactions using differential log records and means for undoing uncommitted transactions using differential log records.

25. (Original) The system of Claim 24, wherein the database comprises a plurality of fixed-size pages.

26. (Previously Presented) The system of Claim 24, further comprising:
means for checkpointing the database by occasionally writing the database in the main memory to one or more persistent backup storage devices.

27. (Currently Amended) The system of Claim 26, wherein the means for checkpointing uses [[the]] a transaction consistent checkpointing policy.

28. (Currently Amended) The system of Claim 26, wherein the means for checkpointing uses [[the]] an action consistent checkpointing policy.

29. (Currently Amended) The system of Claim 26, wherein the means for checkpointing uses [[the]] a fuzzy checkpointing policy.

30. (Currently Amended) The system of Claim 26, wherein the means for recovering replaying comprises:

means for loading the checkpointed database into the main memory database; and
means for loading the log into the main memory database.

31. (Currently Amended) The system of Claim 30, wherein the means for loading the checkpointed database comprises:

means for reading the checkpointed database from one or more persistent backup storage devices; and

means for playing the checkpointed database to restore the main memory database to the state when the backup was made by applying the XOR ~~operation~~ operations between the checkpointed database and the main memory database.

32. (Currently Amended) The system of Claim 30, wherein the means for loading the log comprises:

means for reading the log records from the persistent log storage devices; and means for playing the log records in two [[pass]] passes to restore the main memory database to the latest consistent state.

33. (Previously Presented) The system of Claim 30, wherein the means for loading the log comprises:

means for reading the log records from the persistent log storage devices; and means for playing the log records in one pass to restore the main memory database to the latest consistent state.

34. (Currently Amended) A computer-readable storage medium that contains a program for logging updates and recovering from failure in a transaction-processing system having a main memory for storing a database, one or more persistent backup storage devices for storing a copy of the database in the main memory, and one or more persistent log storage devices for storing log records ~~for parallel logging and parallel recovery~~, where the program under the control of a CPU performs ~~the steps of~~:

generating differential log records by applying bit-wise exclusive-OR (XOR) operations between before-update images and after-update images; and

~~distributing the generated differential log records in parallel to said persistent log storage devices; and~~

replaying the differential log records in an arbitrary order, independent of their generation order, by using the bit-wise XOR operations,

wherein replaying the differential log records further comprises redoing committed transactions using differential log records and undoing uncommitted transactions using differential log records.

35. (Original) The storage medium of Claim 34, wherein the medium is a CD.

36. (Previously Presented) The storage medium of Claim 34, wherein the medium is a magnetic tape.

37. (Previously Presented) The method of Claim 1, further comprising one or more in-memory log buffers wherein each generated log record is temporarily stored in any available log buffer and a group of the buffered log records are written together to an arbitrary one of said one or more persistent log storage devices.